

WHAT IS CLAIMED IS:

1. A system for managing identifiers in a database replication network, comprising:

a database comprising a plurality of data items;

an ID space including a number of identities (IDs) for identifying data items included in the database;

a replica of at least a portion of the data items in the database, the replica comprising an existing range of IDs allocated to the replica from the ID space;

a replica ID manager associated with the replica for requesting a new range of IDs from the ID space when IDs in the existing range of IDs reaches a predetermined threshold; and

an ID administrator associated with the ID space, the ID administrator configured for receiving requests for ranges of IDs, the ID administrator configured for allocating a new range of IDs to the replica in response to the request from the replica ID manager, a size of the new range of IDs being selected based upon an ID usage rate of the replica.

2. The system of claim 1, wherein the ID administrator is a subsystem of the database.

3. The system of claim 2, wherein the database further comprises an interface for receiving the request from the replica

ID manager and for transferring the new range of IDs to the replica ID manager.

4. The system of claim 1, wherein the ID administrator is
5 resident at a different site than the database, and wherein the database and the ID administrator comprise interfaces for communicating with each other via a communications link.

5. The system of claim 1, wherein the replica ID manager is configured for assigning IDs from the replica ID space to data items in the replica to identify the respective data items.

6. The system of claim 1, wherein the ID administrator is configured for allocating the new range of IDs such that the new range of IDs excludes any IDs previously allocated to a replica.

7. The system of claim 1, wherein the ID administrator is configured for calculating the size of the new range of IDs based upon the ID usage rate of the replica.

8. The system of claim 1, wherein the replica ID manager is configured for calculating the size of the new range of IDs based upon the ID usage rate of the replica, and wherein the

replica ID manager includes the size in the request for a new range of IDs.

9. The system of claim 1, wherein the replica ID manager
5 is configured for adjusting the predetermined threshold based upon a usage rate of IDs by the replica.

10. A system for managing identifiers in a database replication network of a database, comprising:

a database comprising a plurality of data items;

an ID space including a number of identities (IDs) for identifying data items included in the database;

a replica of at least a portion of the data items in the database, the replica comprising an existing range of IDs allocated to the replica from the ID space;

an ID manager associated with the replica for monitoring usage of IDs by the replica, the ID manager configured for submitting a request for a new range of IDs, the request comprising a size of the new range of IDs based upon usage of IDs
20 by the replica; and

an ID administrator associated with the ID space, the ID administrator configured for receiving the request from the ID manager and for allocating a new range of IDs to the replica in

response to the request, the new range of IDs comprising the size requested by the ID manager.

11. A method for managing identifiers allocated to one or
5 more replicas of a database, the database comprising a data space comprising a plurality of data items, a global ID space comprising a plurality of identities (IDs) for identifying data items, the method comprising:

allocating a first range of IDs from the global ID space to
a replica of the database;

receiving a request from the replica for a second range of
IDs; and

providing a second range of IDs from the global ID space to
the replica, a size of the second range of IDs being selected
15 based upon a usage rate of IDs by the replica.

12. The method of claim 11, wherein the size of the second
range of IDs is selected based upon at least one of an average
usage rate of IDs, a current usage rate of IDs, and a rate of
20 change of usage rate of IDs over time by the replica.

13. The method of claim 11, wherein the replica comprises a
first replica, and wherein the first replica intermittently

communicates with a second replica of the database for
synchronizing data between the first and second replicas.

14. The method of claim 13, wherein the size of the second
5 range of IDs is selected to provide sufficient numbers of IDs for
the first replica to satisfy ID usage by the first replica
between successive communications with the second replica.

15. The method of claim 14, wherein the second replica
comprises a master copy of the database.

16. The method of claim 11, further comprising submitting a
request from the replica for the second range of IDs, and wherein
the request is received by an ID administrator associated the
15 global ID space, the ID administrator allocating the second range
of IDs from the global ID space in a manner that prevents
conflict with IDs allocated to another replica.

17. The method of claim 16, wherein the request for a
20 second range of IDs is submitted when a predetermined number of
available IDs in the first range of IDs reaches a trigger point.

18. A method for managing identifiers allocated to a
plurality of replicas of a database comprising a data space

including a plurality of data items, and an ID space comprising a plurality of identities (IDs) for identifying data items, the method comprising:

allocating a first range of IDs n_1 from the ID space to a first replica of the database, the IDs in the first range ranging from (x) to $(x+n_1-1)$, x being an integer;

allocating a second range of IDs n_2 from the ID space to a second replica of the database, the IDs in the second range ranging from (y) to $(y+n_2-1)$, y being an integer greater than $(x+n_1-1)$;

receiving a request from a replica of the master database for a third range of IDs; and

providing the third range of IDs from the ID space to the requesting replica, a size n_3 of the third range of IDs being selected based upon a usage rate of IDs by the requesting replica, the IDs in the third range ranging from (z) to $(z+n_3-1)$, z being an integer greater than $(y+n_2-1)$.

19. The method of claim 18, wherein the replica receiving the third range of IDs is the first replica, and wherein the request is generated when a predetermined number IDs from the first range of IDs remain unused by the first replica.

20. A method for managing identifiers allocated to a plurality of replicas of a database comprising a data space including a plurality of data items, and a global ID space comprising a plurality of identities (IDs) for identifying data items, the method comprising:

providing a replica of the database, the replica comprising a replica ID space comprising a plurality of IDs allocated from the global ID space;

monitoring usage of IDs by the replica; and

submitting a request for a new range of IDs from the global ID space when the IDs from the plurality of IDs remaining unused by the replica reaches a predetermined threshold, the request comprising a size of the new range of IDs being requested, the size being based upon a usage rate of IDs by the replica.

21. The method of claim 20, wherein the size of the new range of IDs is selected based upon at least one of an average usage rate of IDs, a current usage rate of IDs, and a rate of change of usage rate of IDs over time by the replica.

22. The method of claim 20, wherein the replica comprises a first replica, and wherein the first replica intermittently communicates with a second replica of the database for synchronizing data between the first and second replicas.

23. The method of claim 22, wherein the size of the new range of IDs is selected to provide sufficient numbers of IDs for the first replica to satisfy ID usage by the first replica between successive communications with the second replica.

24. The method of claim 23, wherein the second replica comprises a master copy of the database.

25. The method of claim 20, wherein the replica intermittently communicates with an ID administrator managing the global ID space, and wherein the size of the new range of IDs is selected to provide sufficient numbers of IDs for the replica to satisfy ID usage by the replica between successive communications with the ID administrator.